

Usable HPC -- An Oxymoron?

Cherri Pancake
Oregon State University
pancake@cs.orst.edu

HPC As a Capital Investment

To increase productivity, substitute capital for labor

HPC requires a significant capital investment -- does it increase productivity ?

Productivity presupposes

- **Making efficient use of expensive resources**
- **Corresponding reduction of human costs**

... How good is HPC's track record?

Oregon State University

HPC Productivity Paradoxes

(1) Ubiquity is no measure of success


It's how HPC machines are used that counts...

Oregon State University

How HPC Machines Are Being Used

Installed base of parallel computers

- **government research centers and national labs**
- **academic institutions and centers**
- **industry R&D sites**
- **commercial and financial institutions**




Key roles

- **(serial) batch server farms**
- **interactive development / testing**
- **parallel production runs**



Oregon State University



Oregon State University

Performance Isn't Enough

Machines are more powerful, but we're using them less efficiently

- 5 years ago, 20% sustained efficiency was respectable
- On today's SMP clusters, it's difficult to get half that

"Effective lifetime" is also a growing concern

- Today's users don't have the luxury of writing to just one machine
- Typical "prime" approximately 2 years
 - » planned obsolescence
 - » Mean-Time-to-Bankruptcy
- Portability is the obvious way to extend software lifetime

Oregon State University

How Portability Affects HPC Costs

Portability isn't just a matter of multiple targets, but multiple moving targets

A ported code only works

- until the new processor boards are installed
- until the shared library changes
- until the next system upgrade
- until system reconfiguration
- until the next reboot
- until the system load changes
- ... or until the next phase of the moon

Oregon State University

HPC Productivity Paradoxes

(2) There's an inherent tension between portability and performance

- “Each new version of each component in the application development environment introduces some new -- though usually justifiable -- quirk” (Mike Frese, Numerex)
- *Experience says:* It's the “portable” code that is most likely to uncover new quirks

Oregon State University

The “Hidden” Cost of HPC

Buying the machine vs. using it for something

- Capital costs vastly out-paced by human costs
- (Some examples)

What parallel applications really cost

- Migrating a parallel application
- Developing a key application
- Best-in-the-business estimate (G. Montry)



“We had this problem with vector computing...”

- J. Worlton's estimate of per-line cost

Oregon State University

HPC Productivity Paradoxes

(3) The more we spend on a machine, the more human effort is required to use it efficiently

Wherever technology is intended to facilitate processes, it's the recurring costs that dominate

- Fixed cost-of-entry = purchase of HPC system and related infrastructure
- Recurring costs = effort expended by humans in order to apply HPC technology

Oregon State University

Usability Is the Real Measure of Success

TRUE or FALSE?
“HPC technology frees us from work”

TECHNOLOGY HAS MET
ITS PROMISE OF REDUCING
OUR WORK LOAD. IT DOES
THIS PRIMARILY BY
PREVENTING US FROM
DOING ANY WORK AT ALL.




Oregon State University

What Usability Is All About

Usability means

- Ease-of-learning
- Ease-of-use
- Usefulness
- Productivity

When will HPC get there?



Oops!

We've been emphasizing machines that perform -- instead of machines that help humans perform better

Oregon State University

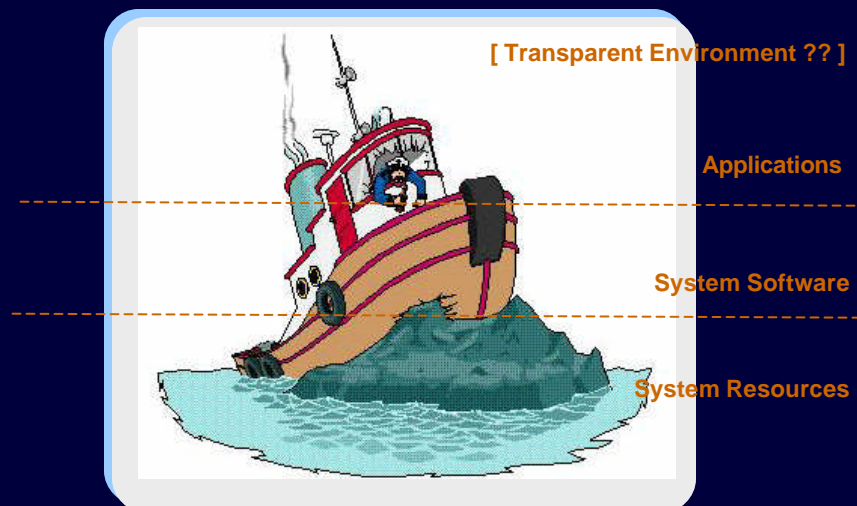
What We Need: *A Healthy Software Infrastructure*



Worlton's "iceberg model"

Oregon State University

What We Have: *Unreliable Software Infrastructure*



Oregon State University

What Application Developers Say about Software

Today's system software and tools ...

- crash!
- don't do what they're supposed to do
- are too machine-specific
- are too diverse and inconsistent
- are not interoperable
- are too complex for users to understand
- don't scale to large applications
- don't scale to large numbers of nodes

Oregon State University

Productivity Paradoxes

(4) Computers in general shift rote tasks to more highly-trained workers

HPC does this very well --

Highly-paid scientists spend ridiculous amounts of time mastering details of technology



Mozart writing the digital version of his symphony No. 38 in D major

Oregon State University

It's Time to Redefine Our Expectations

It's not your father's supercomputer

- **Not as reliable or robust for largescale computation**

The window for success is decreasing

- **Impact of Wall Street**
- **Importance of quarterly performance**

Software capitalization costs are increasing

Oregon State University

What HPC Vendors Say about Software

System software and tool developers ...

- **must scramble just to maintain the status quo**
- **must be able to differentiate their products**
- **get user requests that conflict with company priorities**
- **get user requests that conflict with other user requests**
- **are under-staffed because procurements don't make/break on software**
- **receive mixed messages from users about standards**

Oregon State University

Put Some Muscle behind Software Usability

Involve real users in the procurement process

- **No person in authority understands what it's like to be "in the trenches"**
- **User tasks (not just user codes) need to be part of acceptance**

Develop RFPs that work better for vendors

- **Use standard verbiage so vendors know what to expect**
- **Allow vendors to specify how they fail to meet a specification**
 - » **RFP wording is very "fragile"**
 - » **The discrepancy may not actually be important for your users**
- **Where possible, build in flexibility for delivery dates**
 - » **Does all software have to be there when the machine first arrives?**

Oregon State University

Work to Get More from HPC Vendors

Vendors aren't interested in one-off solutions, even if we're willing to pay

HPC users need to understand the real costs of software

- » **Initial development costs**
- » **Hardening and productizing**
- » **Recurring costs to maintain and support**

Regression testing can soon outpace all development costs!

Form collaborative relationships with vendors


- **Distribution point for software that won't be productized**
- **Provide testbeds to help vendors tackle the market they really want ...**

Oregon State University

HPC Productivity Paradoxes

(5) “Sleeping with the enemy” may be the only way to get usable HPC

We think commercial computing is a threat to technical computing -- but it could be our ticket forward



Natural selection at work

Oregon State University

Learn to Care ... about Commercial Computing

The core business for vendors is mid-range and high-end servers for commercial applications

- Dollar value of this market is at least 6 times that of HPC
- Traditional HPC viewed as a special market with high risks in both R&D and sales

Tools designed for commercial applications may be just what we need

- Compared with large-scale commercial applications, traditional HPC applications are nicely constrained!
- It's much easier to make a case for developing a software tool if it deals effectively with commercial applications


Oregon State University

Is Usable HPC an Oxymoron?

- Emphasizing the wrong things led to HPC that costs more than it yields
- We'll have to work to make that change

The alternative is all too real ...

Large-scale commercial computing



Portability Performance

Produced with (involuntary) artistic assistance from:
Gary Larson, creator of Far Side
Tom Wilson, creator of Ziggy
Scott Adams, creator of Dogbert